Troubleshoot Issues when Element Manager runs in a Standalone Mode

Contents

Introduction

This document provides a summary on how to troubleshoot issues when Element manager runs in a standalone mode.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- StarOs
- Ultra-M base architecture

Components Used

The information in this document is based on Ultra 5.1.x release.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

Ultra-M is a pre-packaged and validated virtualized mobile packet core solution designed to simplify the deployment of VNFs. OpenStack is the Virtualized Infrastructure Manager (VIM) for

Ultra-M and consists of these node types:

- Compute
- Object Storage Disk - Compute (OSD - Compute)
- Controller
- OpenStack Platform - Director (OSPD)

The high-level architecture of Ultra-M and the components involved are depicted in this image:
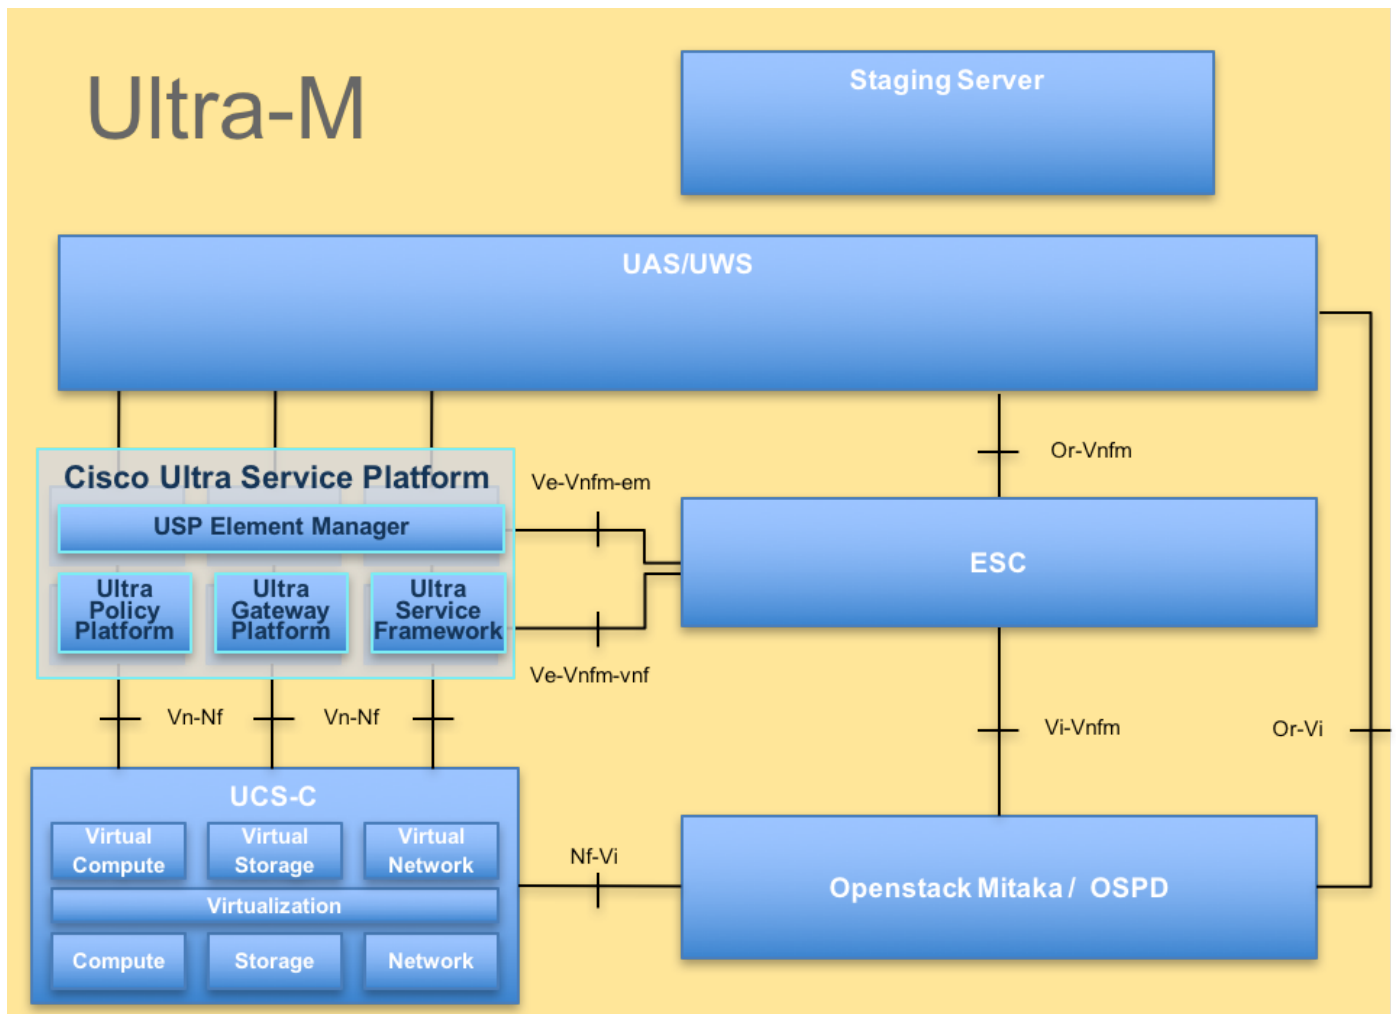
UltraM Architecture

This document is intended for the Cisco personnel familiar with Cisco Ultra-M platform and it details the steps required to be carried out at OpenStack and StarOS VNF level at the time of the Controller Server Replacement.

Abbreviations

These abbreviations are used in this article:

VNF	Virtual Network Function
EM	Element Manager
VIP	Virtual IP address
CLI	Command Line

Problem: EM can End up in this State as it seems from Ultra-M Health Manager

```
EM: 1 is not part of HA-CLUSTER,EM is running in standalone mode
```
It depends upon the version, there can be 2 or 3 EM that runs on the system.

In the case where you have 3 EM deployed, two of them would be functional and third one just to be able to have the Zookeeper cluster. However, it is not used.
In case that one of the 2 functional EMs would not work or is not reachable, the working EM would be in standalone mode.

In case you have deployed 2 EM, in case that one of them is not working or reachable, remaining EM can be in standalone mode.

This document explains what to look if this happens and how to recover.

Troubleshoot and Recovery Steps

Step 1. Verify the State of the EMs.

Connect to the EM VIP and verify indeed the node is in this state:

```
root@em-0:~# ncs_cli -u admin -C
admin connected from 127.0.0.1 using console on em-0
admin@scm# show ems
EM VNFM ID SLA SCM PROXY
3 up down up
admin@scm#
```
So, from here, you can see that there is just one entry in SCM - and that is the entry for our node.

If you manage to connect to the other EM you can see something like:

```
root@em-1# ncs_cli -u admin -C admin connected from 127.0.0.1 using
admin connected from 127.0.0.1 using console on em-1
admin@scm# show ems
% No entries found.
```
Depending on what is the issue on the EM, NCS CLI cannot be accessible, or node can be rebooting.

Step 2. Check the Logs in /var/log/em on the Node that does not Join the Cluster.

Check the logs on the node in the problem state. So, for the sample mentioned, you would navigate em-1 **/var/log/em/zookeeper** logs:

```
...
2018-02-01 09:52:33,591 [myid:4] - INFO [main:QuorumPeerMain@127] - Starting quorum peer
2018-02-01 09:52:33,619 [myid:4] - INFO [main:NIOServerCnxnFactory@89] - binding to port
0.0.0.0/0.0.0.0:2181
2018-02-01 09:52:33,627 [myid:4] - INFO [main:QuorumPeer@1019] - tickTime set to 3000
2018-02-01 09:52:33,628 [myid:4] - INFO [main:QuorumPeer@1039] - minSessionTimeout set to -1
2018-02-01 09:52:33,628 [myid:4] - INFO [main:QuorumPeer@1050] - maxSessionTimeout set to -1
2018-02-01 09:52:33,628 [myid:4] - INFO [main:QuorumPeer@1065] - initLimit set to 5
2018-02-01 09:52:33,641 [myid:4] - INFO [main:FileSnap@83] - Reading snapshot
/var/lib/zookeeper/data/version-2/snapshot.5000000b3
2018-02-01 09:52:33,665 [myid:4] - ERROR [main:QuorumPeer@557] - Unable to load database on disk
java.io.IOException: The current epoch, 5, is older than the last zxid, 25769803777
    at org.apache.zookeeper.server.quorum.QuorumPeer.loadDataBase(QuorumPeer.java:539)
    at org.apache.zookeeper.server.quorum.QuorumPeer.start(QuorumPeer.java:500)
    at org.apache.zookeeper.server.quorum.QuorumPeerMain.runFromConfig(QuorumPeerMain.java:153)
    at org.apache.zookeeper.server.quorum.QuorumPeerMain.initializeAndRun(QuorumPeerMain.java:111)
    at org.apache.zookeeper.server.quorum.QuorumPeerMain.main(QuorumPeerMain.java:78)
2018-02-01 09:52:33,671 [myid:4] - ERROR [main:QuorumPeerMain@89] - Unexpected exception,
exiting abnormally
java.lang.RuntimeException: Unable to run quorum server
    at org.apache.zookeeper.server.quorum.QuorumPeer.loadDataBase(QuorumPeer.java:558)
    at org.apache.zookeeper.server.quorum.QuorumPeer.start(QuorumPeer.java:500)
    at org.apache.zookeeper.server.quorum.QuorumPeerMain.runFromConfig(QuorumPeerMain.java:153)
    at org.apache.zookeeper.server.quorum.QuorumPeerMain.initializeAndRun(QuorumPeerMain.java:111)
    at org.apache.zookeeper.server.quorum.QuorumPeerMain.main(QuorumPeerMain.java:78)
Caused by: java.io.IOException: The current epoch, 5, is older than the last zxid, 25769803777
    at org.apache.zookeeper.server.quorum.QuorumPeer.loadDataBase(QuorumPeer.java:539)
```

Step 3. Verify Snapshot in Question Exist.

Navigate to **/var/lib/zookeeper/data/version-2** and verify that snapshot that is being red in step 2
is present.

```
300000042 log.500000001 snapshot.300000041 snapshot.40000003b
ubuntu@em-1:/var/lib/zookeeper/data/version-2$ ls -la
total 424
drwxrwxr-x 2 zk zk 4096 Jan 30 12:12 .
drwxr-xr-x 3 zk zk 4096 Feb 1 10:33 ..
-rw-rw-r-- 1 zk zk 1 Jan 30 12:12 acceptedEpoch
-rw-rw-r-- 1 zk zk 1 Jan 30 12:09 currentEpoch
-rw-rw-r-- 1 zk zk 1 Jan 30 12:12 currentEpoch.tmp
-rw-rw-r-- 1 zk zk 67108880 Jan 9 20:11 log.300000042
-rw-rw-r-- 1 zk zk 67108880 Jan 30 10:45 log.400000024
-rw-rw-r-- 1 zk zk 67108880 Jan 30 12:09 log.500000001
-rw-rw-r-- 1 zk zk 67108880 Jan 30 12:11 log.5000000b4
-rw-rw-r-- 1 zk zk 69734 Jan 6 05:14 snapshot.300000041
-rw-rw-r-- 1 zk zk 73332 Jan 29 09:21 snapshot.400000023
-rw-rw-r-- 1 zk zk 73877 Jan 30 11:43 snapshot.40000003b
-rw-rw-r-- 1 zk zk 84116 Jan 30 12:09 snapshot.5000000b3 ---> HERE, you see it
ubuntu@em-1:/var/lib/zookeeper/data/version-2$
```

Step 4. Recovery Steps.

1. Enable debug mode so EM stops the reboot.

```
300000042 log.500000001 snapshot.300000041 snapshot.40000003b
ubuntu@em-1:/var/lib/zookeeper/data/version-2$ ls -la
total 424
drwxrwxr-x 2 zk zk 4096 Jan 30 12:12 .
```

```
drwxr-xr-x 3 zk zk 4096 Feb 1 10:33 ..
-rw-rw-r-- 1 zk zk 1 Jan 30 12:12 acceptedEpoch
-rw-rw-r-- 1 zk zk 1 Jan 30 12:09 currentEpoch
-rw-rw-r-- 1 zk zk 1 Jan 30 12:12 currentEpoch.tmp
-rw-rw-r-- 1 zk zk 67108880 Jan 9 20:11 log.300000042
-rw-rw-r-- 1 zk zk 67108880 Jan 30 10:45 log.400000024
-rw-rw-r-- 1 zk zk 67108880 Jan 30 12:09 log.500000001
-rw-rw-r-- 1 zk zk 67108880 Jan 30 12:11 log.5000000b4
-rw-rw-r-- 1 zk zk 69734 Jan 6 05:14 snapshot.300000041
-rw-rw-r-- 1 zk zk 73332 Jan 29 09:21 snapshot.400000023
-rw-rw-r-- 1 zk zk 73877 Jan 30 11:43 snapshot.40000003b
-rw-rw-r-- 1 zk zk 84116 Jan 30 12:09 snapshot.5000000b3 ---> HERE, you see it
ubuntu@em-1:/var/lib/zookeeper/data/version-2$
```

VM reboot might be required once again (would be automatically, you do not need to do anything)

2. Move the zookeeper data.

In the **/var/lib/zookeeper/data** there is the folder called version-2 that has the snapshot of the DB. The error above points the failure to load so that you remove it.

```
ubuntu@em-1:/var/lib/zookeeper/data$ sudo mv version-2 old
 ubuntu@em-1:/var/lib/zookeeper/data$ ls -la
total 20
....
-rw-r--r-- 1 zk   zk      2 Feb  1 10:33 myid
drwxrwxr-x 2 zk   zk   4096 Jan 30 12:12  old  --> so you see now old folder and you do not see
version-2
-rw-rw-r-- 1 zk   zk      4 Feb  1 10:33 zookeeper_server.pid
..
```

3. Reboot the node.

sudo reboot

4. Disable back the debug mode.

sudo reboot

These steps shall bring the service back up on the problem EM.